MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS-1963-A

Color Vision and Computer Vision

Christopher M. Brown
Computer Science Department
University of Rochester
Rochester, NY 14627

SEP 1 6 1983

S

H

AD-A132 520

DTIC FILE COPY

Department of Computer Science
University of Rochester
Rochester, New York 14627

83 09 13 127

Color Vision and Computer Vision

Christopher M. Brown
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 108
June 1982

Abstract

This collection of elementary facts and definitions is meant to be a guide to concepts and results of color vision and color science research that are likely to be of interest to computer visionaries. There are a few thoughts about research topics here, but no results.

Accession For	
NTIS GRA&I	☑
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

1. Levels

The color we perceive an object to have is a fairly reliable indicator of an intrinsic property, the spectral reflectance of its surface: what fraction of impinging light at a given wavelength (or in certain bands of wavelength) it reflects. In a single scene, we get this information via one particular spectral composition of reflected light which clearly depends on the spectral composition of the illumination. One immediate mystery: how is perceived color related to incoming spectral distributions from a scene?

It has long been known that the perceived color of a surface in a complex scene has little to do with the spectral composition of its reflected light (hence the spectral distribution coming into our eyes). Exactly the same spectral composition of light can be experienced as many different surface colors, depending on the properties of spatially and temporally neighboring surfaces. Conversely, different spectral compositions may be perceived as the same surface color (Section 8).

The subjective aspect of color is thus essentially private, and more in the realm of psychology than psychophysics or physics. It is so private and has proved so slippery that in Boynton's *Color Vision* these "subjective phenomena" (i.e., everything we directly experience in everyday life, every color that occurs outside a laboratory) rate 18 out of 437 pages. More on this in Section 9.

Most of this survey has to do with the more manageable aspects of color. We shall see that under sufficiently "reduced" (i.e., artificially simple) conditions, human color perception is rather quantifiable and predictable (Section 5). It is as important to understand basic "color science" as it is to understand that much of it has little to do with everyday color sensations.

Figure 1 illustrates several levels at which it makes sense to describe the visual process. (All the figures appear at the back of the paper.) The physics of light is translated by our neural apparatus into psychophysical responses that are reliably predictable and "well-behaved" in the sense that various useful linear laws hold, color television sets can be designed, and so forth. Psychology is a more private domain, in that it deals with sensations, not (even biological) engineering. Further, many interesting (nonlinear, interdepending) psychological phenomena occur when "real-life" complex scenes are the input to a color visual system. Color TV engineers need not be concerned with them, since they count on a human visual system to provide them. Computer visionaries and cognitive scientists may choose to take up the challenge of recreating, modelling, explaining, or implementing them.

Sections 2 and 3 outline concepts from physics and psychophysics. This paper would be a good spot to outline the facts on the primate visual system anatomy, but I shall give that relatively short shrift here (Section 4). Of course ultimately one wants to understand neural function, and to make neural models. Another issue that is not addressed here but seems important and illuminating is abnormal color perception.

2. Physics of Light

The physics of Newton's *Opticks* (1730) is good enough for us. We just need to know that visible light occupies the wavelengths (approx.) 380-750 nanometers of the electromagnetic spectrum. Sunlight spans this range (sunlight through the atmosphere varies from about 200 to well above 1000 nm.) We are sensitive to those wavelengths that are most strongly represented in sunlight filtered through the atmosphere.

Colors can arise from interference and diffraction phenomena (some butterfly wings, some bird feathers, some oil slicks, some glitzy jewelry, etc.); I won't mention that further. They also arise in temporal phenomena--certain structured black and white flickers (e.g., Benham's top) give rise to the perception of color: enough said on that, too. The color I mainly consider in this paper arises from mixtures of light. In everyday life, most colors arise from light reflecting off pigmented surfaces. The pigments absorb wavelengths differentially, hence modulating the spectral content of reflected light (Figure 2). For technical reasons we usually speak in terms of light mixture rather than absorbancy mixture.

The perception of achromatic (non-color) luminance (loosely, brightness) closely obeys a law of superposition or addition called *Abney's Law* (see Section 3.2). Thus by definition the perceived luminance of a "source" is the integral of its radiant energy in different wavelengths times the sensitivity of our visual system in different wavelengths. Colors behave differently; they do not merely add, but they can cancel each other as well, producing (for light mixtures) colorless (white) light. In certain experimental situations, the color vision system behaves according to the empirical *Grassman's Laws*, which describe the linear, signed nature of color mixture. Unlike our auditory system, our visual system does not do "Fourier analysis" on light--we cannot identify the components of a color mixture from the mixture. This last statement is not obvious, but see Section 5. The issue of the visual system doing spatial frequency decomposition is a completely different one.

3. Radiant Energy, Luminosity, and Color

This section presents in a rather bottom-up order some basic definitions that go with the physical, psychophysical, and psychological levels. It draws on [Sheppard 1968, Horn and Sjoberg 1978].

3.1 Radiant Energy: Physics

Radiant Flux P: refers to the rate of energy U being emitted, transferred, or received in the form of radiation.

$$P = dU/dt \qquad \text{(watts)}$$

Radiant Intensity J: refers to a source of radiant energy, and is the radiant flux emitted per unit solid angle in a given direction θ. (There are 4π *steradians* in the sphere of all possible directions). The total flux of a source is the integral of J over all directions.

$$J(\theta) = dP(\theta)/d\omega \qquad \text{(watts/steradian)}$$

Irradiance H: refers to a surface receiving radiant energy. It is the incident radiant flux per unit area of surface (incident flux density).

$$H = dP/dA \qquad \text{(watts/sq. meter)}$$

Radiant Exitance M: refers to a surface emitting radiant energy. It is the exitant radiant flux per unit area. Total radiant exitance equals total irradiance if the surface reflects all incident light, transmitting and absorbing none.

$$M = dP/dA \qquad \text{(watts/sq. meter)}$$

Radiance N: refers to a source and is the flux emitted per unit foreshortened surface area per unit solid angle. If the angle between the surface normal and direction of exitant radiation is θ, then the foreshortened area is the actual surface area times the cosine of θ. Radiance is equivalently the flux emitted per unit surface area per unit projected solid angle. The apparent Lightness (Section 3.3) of a surface patch is related to its radiance. *Image irradiance* is the irradiance through an optical system that falls on the image plane. *Scene radiance* is the radiance of an imaged scene, and in usual imaging systems is proportional to the image irradiance [Horn and Sjoberg 1978].

$$N(\theta) = dJ(\theta)/dA(\theta)$$

$$N = d^2P/((dA \cos\theta \, d\omega) \qquad \text{(watts/((sq. meter) (steradian)))}$$

3.2 Luminosity: Psychophysics

Only certain radiant energy elicits visual response (~380-570 nm). Response varies nonlinearly with wavelength, holding energy constant (Figure 3). The reaction of an ideal, "average" eye is given by standard Luminous Efficiency Functions $V(\lambda)$ and $V'(\lambda)$, published in 1924 by the CIE *(Commission International de l'Eclairage)*. The functions vary between 0 and 1 in dimensionless units. V' is a luminous efficiency function for low light levels (only rods respond--Scotopic vision) and V is for high light levels (rods and cones respond--Photopic vision).

For each wavelength of radiation λ, each radiant quantity above has a corresponding photopic quantity (response it elicits) that is just the product of $V(\lambda)$ with the quantity. It is assumed that luminances for different wavelengths add linearly. The expression of this assumption is *Abney's Law*, that says the *integral* of this product over wavelength is the total luminance of a source (see Luminous Flux below).

To make things complicated, a different "energy" unit for luminous quantities is chosen: the unit is the *lumen*, and there are K = 678.8 lumens per watt. A lumen is a *candelasteradian*, and a candela is the unit of photopic luminous intensity I. The *candela* is chosen so that the photopic luminance L of a full (black body) radiator at the temperature of freezing platinum is 60 candelas/(sq. cm.). So a lumen is the flux emitted in a unit solid angle (1 steradian) by a point source having a uniform intensity of 1 candela. A *foot-candle* is the unit of illuminance E: it is the illuminance provided by a point source of 1 candela intensity at the distance of one

foot: it is 1 lumen/(sq. foot). Last, a *lambert* is the uniform luminance of a perfectly diffusing surface emitting light at the rate of 1 lumen/(sq. cm.). Luminous units are thus $KV(\lambda)$ times radiant units.

Luminous Flux F (photopic version of radiant flux): P is in watts.

$$F(\lambda) = KV(\lambda)P(\lambda)$$

$$F = K \int V(\lambda) P(\lambda) \, d\lambda. \qquad \text{(lumens)}$$

Luminous Intensity I (photopic version of radiant intensity).

$$I(\theta) = dF(\theta)/d\omega \qquad \text{(candelas)}$$

Luminance L (photopic version of radiance).

$$L(\theta) = dI(\theta)/dA(\theta) \qquad \text{(candelas/(sq. cm.)) or (stilb)}$$

Illuminance (or *Illumination*) E (photopic version of irradiance).

$$E = dF/dA \qquad \text{lumen/(sq. cm.), alias (phot)}$$
$$\text{or lumen/(sq. meter), alias (lux)}$$

Retinal Illumance (unit is the *Troland*): A troland is the illuminance produced by viewing a surface of luminance 1 candela/(sq. meter) through a pupil of 1 sq. mm. area. Used in reporting psychophysical experiments.

Last, an important point. The basic idea in these psychophysical units is to be able to predict the effect of physically measurable radiation. The above units are based on some statistical average observer, and are (if you like) for achromatic colors. Very much the same exercise is carried out for chromatic colors, and produces the CIE "standard [color] observer" response curves, and various co-ordinates for color response space.

3.3 Color: Psychology

Aperture color, reduced color: color such as that experienced as filling a hole in a screen (usually small portion of visual field). Not concentrated in a plane or spread throughout a volume, nor located in depth. Attributes of hue, saturation, brightness. (Also chromaticity coordinates of various sorts).

Surface color (sometimes *object color*): color experienced as property of a surface. These are the colors we usually perceive. They are "hard," they "resist the gaze," and they inherit the plane and texture of their surface. Attributes of hue, saturation, and lightness.

Volume color: color experienced as property of bulk material.

Memory color: a remembered appearance--can influence perceived color.

Hue: color attribute that permits color to be classed as red, yellow, green, blue, and their intermediaries.

Chromatic colors: colors with hue.

Achromatic colors: colors without hue. Achromatic aperture colors range from very dim to very bright. Achromatic surface colors range from black through gray to white.

Saturation: the attribute of a chromatic color that determines the degree of its difference from the achromatic color most closely resembling it. Measures the "lack of white" mixed with the color.

Brightness: the attribute of any aperture color that permits it to be classed as equivalent to some achromatic color ranging from very dim to very bright.

Lightness: attribute of any surface color that permits it to be classed as equivalent to some member of the series of grays from white to black.

4. Neural Processes

The human visual pathways are shown in no detail in Figure 4. Light is focussed on the retina, where it excites rods containing the pigment rhodopsin and cones containing three sorts of pigment. Cones containing erythrolabe are sensitive to light in the longer wavelengths; cones with chlorolabe are sensitive to middle wavelengths, and cones with cyanolabe respond best to shorter wavelengths. Cones are most populous in the fovea, the high-resolution spot on the eye that is the center of our visual field. Cone population density drops off toward the periphery of the retina. Rods are more sensitive than cones, and respond (in scotopic vision) to light levels below the cone threshold. Moonlight seems to provide scotopic conditions. Photopic vision begins at about twilight conditions as the cones come into play. The long-wavelength pigment is the most sensitive and the short-wavelength pigment is least sensitive. Red is the first color we become aware of as illumination levels rise.

You might expect individual responses to light to vary, and indeed they do (this seems to be an active area of work currently in color perceptual circles). However, an average curve of photopic spectral sensitivity (that will keep coming back) is shown in Figure 3. This function is known as V, and there is a similar one, shifted a bit left, called V'--it describes scotopic sensitivity. It is a considerable technical exercise to get at the spectral absorbancies of the three cone pigments and to factor in the absorbancies of the biological structures; however, psychophysical response curves for the three pigments do exist. Their sensitivity curves are shown in Figure 5, but note that the curves have been shifted up and down relative to one another for the purpose of computing opponent sensitivities (see next paragraph). In the curves of Smith and Pokorny giving the sensitivies of the pigments, the short-wavelength pigment is almost 100 times less sensitive than the red.

No one really seems to know what computations are performed in the primate retina. It is a bona fide extension of the brain, connected by the optic nerve. There are several sorts of neurons in the retina connected to the rods, cones, and each other in varying topologies (it seems). It has been suggested that the retina could do "lightness" calculations to produce a response that is sensitive to reflectance (an intrinsic property of surfaces) rather than apparent brightness (that varies with

illumination) [Marr 1974]. It has been suggested that the retina could perform a transformation from the basically "red, green, blue" parameterization of light suggested by the pigments to an "opponent color" space more in line with psychological facts and anatomical findings [DeValois and DeValois 1975] (Figures 5 and 6). The psychological fact is that we do not have names or perceptions that go with the concepts "reddish-green," "bluish-yellow," or "whitish black." There is also some neurophysiological evidence that there are cells early in the system (perhaps retinal ganglion cells, certainly lateral geniculate nucleus cells) that respond to differences of opponent colors (Figures 6 and 7). There is more evidence that dilutes the foregoing evidence somewhat, but there does seem to be reliable psychophysical evidence that at some level there are privileged "channels" of Red-Green and Blue-Yellow that act independent and orthogonal [Williams et al. 1980].

Many neurophysiological measurements are made in the LGN, which is usually taken to reveal more or less what goes on in retinal ganglion cells. The LGN has a pretty layered structure in which alternate layers receive input from ipselateral and contralateral foveas; the layered structure washes out in the LGN for input from the periphery.

The striate cortex, where the visual pathway leads next, is pretty much a mystery as far as color is concerned, but there are some studies on monkeys [Zeki 1980]. They may indicate that color perception is intimately associated with shape perception by then.

Abnormal color vision takes several forms; monochromats are not very revealing. There are several flavors of dichromats: Protanopic, Deuteranopic, and Tritanopic. These flavors arise from the lack, substitution, or anomaly of one or another visual pigment. Normal dichromatic perception is relatively well explained as pigment loss. The perceivable (aperture) colors are those you would expect if "normal" three-dimensional (R,G,B) colors were projected onto the two-dimensional space spanned by two pigments. The result is prettily pictured as lines of confusion on a chromaticity diagram (see below). There are several more or less puzzling variations on this "simple" dichromatism. For more on all this, see [Boynton 1979].

5. Perception of Aperture Colors

Aperture color perception is the basis for *colorimetry*, which within its limitations is a pretty and well worked-out descriptive theory. It neither says nor purports to say anything about sensations. In the experiments that establish the basic data, human subjects serve strictly as nulling devices. The main experimental paradigm is that of color matching; half an optical field contains an aperture color to be matched, and the subject manipulates controls until the other half the field matches the test half. No reports of color names or sensations are required: only the matching behavior.

It is found under these conditions that three variables are necessary and sufficient to match any presented color. This is not entirely surprising, given the three degrees of freedom furnished by the retinal pigments, but it is nonetheless gratifying that every color can be matched using only three degrees of freedom. These experiments discover *metameric* matches. Metamerically matched (indistinguishable perceptually) colors may arise from wildly different spectral compositions. Two colors arising from identical spectral distributions are called *isomeric*.

The informal results of matching experiments are the following. For any monochromatic light of wavelength $\lambda 1$ less than about 490 nm, there's another $\lambda 2$ greater than about 570 nm that can mix with $\lambda 1$ to form white. If we add two lights not satisfying those wavelength restrictions, EITHER there is a third monochromatic light of wavelength $\lambda 3$ that, when added to white, will match the mixture of $\lambda 1$ and $\lambda 2$ OR there is a monochromatic light of wavelength $\lambda 4$ that, when added to $\lambda 1$ and $\lambda 2$, matches white. This last is true when $\lambda 1$ mixed with $\lambda 2$ is a purple ($\lambda 1$ and $\lambda 2$ are reddest and bluest spectral colors); then $\lambda 4$ is greenish. These observations give three *metameric equations* (which are not really irreducible). Below, $L1 - Ln$ are luminances, $\lambda 1 - \lambda n$ are wavelengths. W is white light, $++$ means "mixed with," $==$ means "matches."

(1) $L1(\lambda 1) ++ L2(\lambda 2) == Lw(W).$

(2) $L1(\lambda 1) ++ L2(\lambda 2) == Lw(W) + L3(\lambda 3)$

(3) $L1(\lambda 1) ++ L2(\lambda 2) + L4(\lambda 4) == Lw(W)$

"Grassman's Laws" say essentially that the above matching equations may be treated like equations involving real numbers. The mixing operator $++$ corresponds to addition or $+$, $==$ is transitive, $==$'s may be substituted for $==$'s, $==$ is preserved by addition or subtraction of $==$ terms on both sides, both sides of an $==$ may be multiplied by a constant. Note that Abney's law says luminances are additive, but Grassman's laws say colors can neutralize each other. From this it follows by induction that we can reduce any number of monochromatic luminances (and in the limit *any continuous spectrum*) to a mixture of white light and a single colored luminance, which is either monochromatic or a *"pure purple"* (a mixture of "pure red" and "pure blue").

(4) $\Sigma_i \; Li(\lambda i) == Lw + L\lambda$

Note that Equation (4) has three explicit variables: two luminances Lw and $L\lambda$, and λ indicating something like "hue" or "wavelength or purple." The infinity of variables required to specify complex radiation in physical terms is reduced to three

independent variables for matching. This is the *trivariance of color vision*. Note again: Matching has nothing to do with sensation.

Colorimetry assumes trivariance of vision, and has defined a *standard observer* that will always give the same response to a stimulus. First, convert the above equations into standard algebraic ones and write

(5) $\Lambda = L1(\lambda i) + L2(\lambda 2) + L3(\lambda 3)$.

Λ is any stimulus, $\lambda 1 - \lambda 3$ are a *primary system*--3 suitably chosen monochromatic stimuli called "primaries." Equation (5) may not directly express a match, since L's may have to be negative (if Λ is a monochromatic light, for instance.) Negative luminances are nothing to fear--they just mean their terms were on the opposite side of the metameric matching equation.

A primary system must provide a basis for a color space. Thus it must have no primary match a combination of two others, and some combination of primaries must give white light. There are an infinity of primary systems, interconvertible through linear transformations. Given a primary system of λi's, (L1, L2, L3) specify Λ unambiguously. Wright used primaries of 460, 530, and 650 nm to get his curves (Figure 8). They display fairly intuitive results except that some coordinates are negative, indicating negative luminances.

Wright's data was smoothed up and became a CIE standard (the *RGB system*). It has the disadvantage that at least one coordinate is negative for monochromatic light, and CIE was afraid the arbitrarily chosen primaries might inherit spurious physiological merit. So we take the next step.

Assume different luminous units P1, P2, P3 are chosen for the three primaries $\lambda 1$, $\lambda 2$, and $\lambda 3$ in Equation (5). Then rewrite (5) as

(6) $\Lambda = G1(P1) + G2(P2) + G3(P3)$

where the Gi are amounts of the (now unequal) unit quantities Pi. Trichromatic coordinates are defined

(7) $gi = Gi / (G1 + G2 + G3)$ $(i = 1,2,3)$

so

(8) $g1 + g2 + g3 = 1$.

This amounts to projecting the Gi coordinates onto the "unit plane" defined by Equation (8). Define a new luminous unit for Λ as

(9) $\Lambda = 1/(G1 + G2 + G3)$;

then Equation (6) may be written

(10) $\Lambda = g1(P1) + g2(P2) + g3(P3)$.

where the left hand side is one new luminous unit of light Λ. So (g1, g2, g3) give the fraction that each primary contributes to the total (taken as unity) in matching the

given stimulus. The g_i thus emphasize the chromatic aspects, not the luminous aspect, of the stimulus.

The CIE *International XYZ System* has three purely imaginary primaries such that: (A) The chromaticity coordinates for any real stimulus are all nonnegative. (B) Luminous units of zero are chosen for primaries X and Z (!!). The luminous unit for Y is unity--the Y coordinate accounts for all the luminosity in the stimulus. So Y is a good thing to send to your black and white TV set. (C) The last feature implies the need for one more condition to relate XYZ to RGB: that "equal energy white" (EEW) matches the mixture with equal amounts of the X, Y, and Z primaries. Equal energy white has uniform radiant flux at all wavelengths.

The color-mixture properties of the CIE 1931 Standard Observer are given in the curves for x, y, and z (Figure 9). These *distribution coefficients* are the tristimulus values of the spectrum for the Standard Observer. By feature (B) above, the y function is the same as the 1924 CIE Photopic Luminous Efficiency Curve $V(\lambda)$ that tells how a standard observer reacts to light of a given wavelength and unit energy (Figure 3). For any source of spectral radiance P (λ),

$$(11) \quad X = \int P(\lambda) \; x(\lambda) \; d\lambda$$

$$Y = \int P(\lambda) \; y(\lambda) \; d\lambda$$

$$Z = \int P(\lambda) \; y(\lambda) \; d\lambda$$

The chromaticity coordinates x, y, z are, in analogy with Equation (7),

$$(12) \quad x = X \; / \; (X + Y + Z)$$

$$y = Y \; / \; (X + Y + Z)$$

$$z = Z \; / \; (X + Y + Z)$$

The (x,y) chromaticity diagram results from projecting away the z coordinate as shown in Figure 10. A color is a point in the chromaticity diagram. Its hue is often defined as its *dominant wavelength*, which is the point on the boundary of the chromaticity diagram intersected by a half-line from EEW through the color. The boundary of the diagram contains the pure (completely saturated) spectral colors and the "pure purples" (mixtures of pure red and pure blue). The *saturation* of a color can be measured by its distance from the boundary in the direction of the (completely desaturated) EEW.

The 1931 CIE chromaticity diagram was made for stimuli of 2-degrees angular subtense (Figures 11, 12). A 10 degree one also exists, and is subtly different. Many rescalings of chromaticity diagrams are also useful, e.g., to provide uniform perceptual scaling [Wandell 1982]. The CIE also specified four spectral distributions for illumination, all easy to obtain with a tungsten light and filters. Illuminant A is like indoor light, B is like noon sun, C is like normal daylight, and D6500 is even more like normal daylight. Equal energy white is a theoretically important illuminant that is hard to obtain in practice. An interesting line in a chromaticity diagram is the chromaticity of black body radiation; (dark through deep red, orange, yellow, white, blue-white). All the standard illuminants are close to it.

Again, it is important to emphasize that the CIE system is based on an artificial observer who responds in just one way (by matching) to chromaticity and brightness. Sensation is not involved, nor is the subjective appearance of colors. A fortiori the appearance of colored surfaces in complex spatial scenes is not addressed. There is evidence by Evans that for scenes with light sources, illumination, reflectance and transmission, four variables are necessary to describe color sensation. Also there is convincing evidence (see Section 7) that a surprising range of surface hues can be evoked from bivariant stimuli. The visual system has the power to perceive almost a full gamut of hues when presented with mixtures of only two narrow spectral bands: this is the "Land effect." In any event, the CIE never had pretentions to a "theory of vision," and it is a mistake to refer to various surface color phenomena as "departures from classical expectations" because they are not predictable from classical colorimetry.

6. Color spaces

Color spaces that account for chromaticity and brightness, even in simple matching experiments, must be three-dimensional. Leaving out the possibility of higher-dimensional spaces, there has been much work in 3-D color spaces for various purposes. This little catalog is merely to hint at this diversity [Meyer and Greenberg 1980].

The CIE *XYZ* space is the subject of Section 5.

The *RGB* space for a given set of primaries of three particular chromaticities is interconvertable with the XYZ space by a simple linear transform: the origin moves to the desired achromatic point, and the axes rotate and skew to go through the primary chromaticities. RGB is the natural space when you are sending colors to an RGB monitor.

Opponents space is good for naming colors (Figure 13); basically its axes are White-Black, Red-Green, and Blue-Yellow. This space was discussed in Section 4.

The CIE has defined many other spaces, some scaled so that equal distance in the space corresponds to equal perceptual differences [Meyer and Greenberg 1980].

A broadcast TV monitor decodes the *YIQ* color space. Y is the familiar CIE Y, which encodes all the luminosity information. A Black and White set uses this and throws away the rest. I and Q are again derived by a linear transformation of X, Y, and Z, chosen for reasons not of interest here.

Munsell published a book of color chips and a nomenclature based on a hue, value, chroma system (see [Judd and Wysecki 1963]). Hue here is defined in terms of chromaticity and luminosity--this gives a little freedom to represent more than pure CIE chromaticity. For instance, there is no "brown" in CIE chromaticity. Value is an expression of luminosity, basically CIE Y. Chroma is like saturation, measuring departure from the nearest achromatic color. The relation of Munsell co-ordinates to XYZ can be made precise, but it is not a linear transformation.

HSV space (hue, saturation, value) is like Munsell space. It is more natural for artists or computer graphics people who think about mixing paints.

7. Phenomena of Adaptation and Spatial Interaction

Helmholtz (1853) knew of most of the failings of the above descriptive theory. The relevant effects also have parallels in surface color perception. These important effects hinge upon interaction between colors, spatially or temporally. This section draws on [Boynton 1979 (Chapter 2)] and [Sheppard 1968].

Brightness Adaptation. A central thing to remember about the visual system is that it is constantly adjusting the "zero point" of its instantaneous achromatic and chromatic perceptual range. The dynamic range of the system is enormous, but not all of it is available at any one time. Photopic vision is effective over a range of brightness (varying by a factor of a million or more) much larger than can be explained by changes in pupil diameter. Bleaching of pigment (or other peripheral adaptation) may account for some of the remainder. Central mechanisms may also be involved.

Chromatic Adaptation. There are chromatic adaptation effects that take place with timecourses from milliseconds to minutes. The system seems to adapt to ambient (or average) illumination, for instance, in non-central ways.

Achromatic Surrounds. Perception of whiteness of a spot depends on its surround--a darker surround brightens the spot, a brighter one increases the "gray content" of the spot. A bright enough surround forces an initially white spot to be perceived as black. A spot by itself will never be perceived as black, no matter how dim. This effect could be explained by brightness adaptation and saturation of the low and high ends of the instantaneous dynamic range. With a chromatic spot and achromatic surround, a low relative luminance of the colored spot to the surround leads to a perception of the spot as black. As the ratio of brightnesses increases, increasingly lighter gray seems to be mixed with the color. The grey disappears at a luminance (sometimes very much) below that of a luminous match with the surround. Next the color seems to glow and its saturation (lack of admixed gray) continues to increase. When the color is slightly brighter than the surround, the color begins desaturating and taking on the appearance of a light source. The Gibbs effect is that a brightly illuminated object, even a piece of coal, can be perceived as white if the light source and beam are undetectable. A highly reflective object thrust into the beam causes the object to be perceived as dark.

Dark Colors. Brown is an undeniably experienced hue. Brown is produced by orange or yellow with a bright surround. Dark colors (especially black and brown) rarely appear without a surround. Other dark colors (e.g., maroon, olive drab) seem related to spectral hues, but all have in common the need for a surround.

Lighting and Surface Effects. Perceived colors are affected by the perceived shininess of surfaces, which usually correlates with "hardness." Colors on surfaces with highlights tend to look more saturated than colors on the same surfaces diffusely lighted (no highlights, less contrast). It is worth mentioning that highlights give more information about the chromaticity of the illumination than about the

reflectance of the surface [Cook and Torrence 1981].

Chromatic Surrounds. Simultaneous Chromatic Contrast is the phenomenon that a colored surround in general affects the perception of a colored spot. Changes in lightness and hue occur; contrast effects are by definition of a "complementary" flavor (the surrounded color is driven "away" from the surrounding one). The so-called "spreading effect" directly contradicts this, however: under some conditions colors "co-operate" and dark surrounds darken the surrounded color. In the colored shadow effect, a desaturated wash of color is preceived as white, but a shadow cast in it has a pronounced complementary hue.

Afterimages and Successive Contrast. Afterimages of complementary colors are easily produced on the retina. The perception of a color A can be violently affected by fixating on a different color B first.

8. Surface Colors

It has long been known that the subjective aspect of surface color perception is not related in any obvious way to aperture color perception. One of the most important phenomena in surface color perception is called *color constancy*. Here I want to hint at the effect and make three pointers into the literature.

Color constancy is the effect that perceived colors in real-world scenes remain stable over a wide variety of illumination spectral compositions. In computer visionese, humans can compute intrinsic color images. It has been convincingly demonstrated that as its spatial color context changes, we can see the same spectral distribution (i.e., isomeric, a fortiori metameric, aperture color) as a range of colors that spans the entire range of hues. This ability specializes to achromatic colors, for which it is called "lightness constancy" [Marr 1974, Horn 1974].

Color constancy effects are so undeniable and strong that some have been led to deny any connection between the perception of surface colors in a real scene and aperture color perception. This is probably too strong a holistic stance, but the evidence is relatively strong as well. There are two ill-differentiated approaches to constancy that I have been able to find. One was enunciated by Helmholtz (and probably before). It has come to be called "discounting the illuminant," a special case of the "taking-into-account" school of constancy that basically wants to subtract the effects of eye movement to get position constancy, illumination to get color or brightness constancy, etc. Helmholtz would have spoken of "unconscious inference" to explain constancy--that does not seem necessary, though some are still fond of it [Rock 1977]. The other main school seems to be a "higher-order cue" school, who maintain that the visual system uses relational information in the image to compute really relevant (more or less "intrinsic") images. This may have begun with the Gestaltists, and certainly J.J. Gibson [1950] is a visible figure of this school. This school has had trouble coming up with reasonable mechanisms for computing intrinsic images, and many psychologists as a result tend not to treat them seriously.

Judd in 1940 [Judd 1979] gave an explanation of what it means to discount the illuminant in surface color perception. It was backed up with his usual careful experimental work. The basic idea is very simple; somehow compute an average

spectral radiance of the complex scene. If there is a wide diversity of surface colors in the scene, the average is approximately what a gray surface would reflect under the illumination, and hence is closely related to the illumination itself. Define the chromaticity of this average to be a new origin in the chromaticity diagram. Then the hue of any surface color in the scene is given by the *direction* of the vector from the new origin to its chromaticity in the chromaticity diagram. Its saturation is determined by the *length* of the vector--the longer, the more saturated. The normal dominant wavelength definition of hue is a special case of this definition, with the illuminant at equal energy white. The precise position of the new origin is calculated using several second-order terms arising from more sophisticated weightings of surrounding colors based on data from Helson.

Land (about 1975) [Land 1959, 1977] devised a theory of surface color perception that can easily be explained as the invention of a new color space. The new space is parameterized by three quantities that are basically relative lightness in three spectral bands (corresponding to spectral sensitivity bands of cone pigments). The lightness is relative to that (in the given band) which is expected of a white surface under the given illumination. There are many possible mechanisms for computing relative lightness. Land proposed a mechanism based on the fact that the eye seems to respond to edges, and to respond proportionally to the ratio of brightness change across an edge. By travelling in a path from region A to B, successively multiplying the ratios across brightness edges, the ratio of A's to B's brightness in any given band may be computed. The cumulative ratio is simply normalized to 1 whenever it goes over 1; after several redundant paths have been computed, the brightest region (in each color band) has been found, and all ratios are correct. The set of three ratios (suitably normalized) provides a consistent explanation of perceived color. Slowly-varying illumination may be discounted by ignoring ratios near 1 along the path. Land's earlier work in two-primary projection demonstrated convincingly that the visual system is sensitive to coherently varying ratios of intensity in two spectral bands, which it perceives as almost a full gamut of chromaticities, rather than the simple mixture of the two bands that would occur in aperture color perception.

Judd's theory, basically one of adaptation, seems likely to get into trouble when illumination varies slowly over the scene, or when the average reflected spectral intensity does not approximate the illuminant's spectral intensity. Of course people may get in trouble there, too. Judd gives a reasonable re-explanation of Land's two-color projection data, but fails to account in a convincing way for the conditions under which such displays look achromatic [Sheppard 1968], or with the Land experiment which obtained a wide gamut of perceived colors by projecting two very close wavelengths of yellow light. Judd's quantitative model has the flavor of being a polynomial fit to several parameters, with the resulting predictable instabilities (see [Horn 1974]).

Horn [1974] demonstrates a method for computing relative ratios that is different from Land's, but one that he specifically relates to Land's work. Horn applies his method to achromatic color constancy (computing lightness), but Land predicts that a threefold repetition of the exercise yields color. Broadly, Horn considers the visual system to respond to the Laplacian of brightness variation, he gets an analytic inverse for the Laplacian, and implements that to extract reflectance (i.e., lightness) from observed intensity images. Low spatial frequency brightness effects presumably due

to illumination changes can be removed by thresholding. Marr showed how something like this process could be implemented in the primate retina [Marr 1974].

None of these three "theories" of color constancy considers surfaces at different slants. Gilchrist has shown [Gilchrist 1979] that perceived slant can affect perceived lightness. His claim is basically that suface colors perceived to be coplanar exert stronger influences on one another than those perceived to be non-coplanar, even keeping retinal adjacency of the regions constant. So exactly the same retinal regions cleverly distributed on a clever 3-D object can be perceived to have lightnesses that differ violently as viewing changes from monocular to binocular. It would be interesting to look into analogous situations in color perception [Brown 1982].

9. Computer Color Vision

Much of recent computer vision work attempts to extract intrinsic images, or physical scene parameters, from digitized images. In an achromatic image, usually the scene radiance is represented directly, having been captured by some imaging device. Radiance, even in complex scenes, is closely related to our private perception of luminance, or brightness, and so it seems natural to talk about "computer vision"-- our programs have a comparable input to our sensations. A chromatic image is usually represented by a set of "overlay" images, giving scene radiance in several spectral bands. This is not what we "see"--unlike the auditory system, the visual system does not allow us to analyze spectrally complex inputs into frequency components. We are aware of (subjective versions of) hues, saturations, lightnesses, but we know very well that we cannot derive spectral composition from them.

Thus the very input to color computer vision contains an "inhumanly" high amount of intrinsic information. One obvious track then would be just to do "lightness" computations in however many spectral bands are represented in the input. Insofar as this is successful, it may derive information about the illuminant or about the reflectance properties of the surfaces. Note that this approach (as does most of "color science") leaves out all subjective aspects of color and might just as well be called multi-spectral lightness computation. Such a program can be carried out purely on the basis of physical principles, and need not refer to human anatomy, behavior, or senses at all [Horn 1974]. Psychological effects such as "color constancy" would follow from accurate reflectance computations. There are indeed some unsolved problems in brightness calculation.

- (Achromatic) Reflectivity calculations have traditionally been carried out under assumptions of slowly (spatially) varying illumination, quickly varying reflectance, and planar surfaces. Obvious questions: what about quickly varying illumination, slowly varying reflectance, or non-planar surfaces?

- Shape from (achromatic) shading calculations have traditionally been carried out under assumption of spatially non-varying (and known) illumination, fixed reflectance, and nonplanar surfaces. Similar obvious questions.

Thus there is ample room for introducing increased complications into the calculations of these intrinsic images [e.g., Brown, Ballard and Kimball 1982]. The mechanisms of interaction between the various intrinsic images are of course the technical focus of any such work.

Recent work at MIT [Rubin and Richards 1981] investigates how color (in the form of spectral samples of an image) can indicate possible changes of material in the scene (as opposed to changes caused by shadows, varying pigment density, shading variations from varying orientations, or highlights).

Another interesting research problem is to predict or explain human color matching, color naming, or color sensation (e.g., constancy) in complex surface color scenes. This requires a theory of human color processes. I see Land's theory of human surface color perception as basically being a luminosity calculation (a lightness calculation that takes into account certain human parameters (such as the spectral absorption of eye pigments). I suspect that such a pure approach is not adequate, and in fact I have no feeling about how heavily it should be weighted in comparision to, nor how it interacts with, various (spatially and temporally) global and local adaptation and contrast phenomena. Not to mention its interaction with surface slant data, which seem to affect lightness calculations. Let alone its interaction with "memory colors" and other more or less cognitive processes. A daunting approach to this problem would be simply to integrate and weigh a large number of interdependent but individually understood effects with no theoretical foundation. Even that would probably be new. Much better would be a model from which fell a satisfying diversity of all these effects.

References

Boynton, R.M. *Human Color Vision*. New York: Holt, Rinehart and Winston, 1979. Standard text on color vision.

Brown, C.M. LandLab: Experiments in computer graphics and color vision. Technical Report in preparation, Computer Science Dept., U. Rochester, expected 1982.

Brown, C.M., D.H. Ballard, and O. Kimball. Shape and light source direction from shading. Technical Report in preparation, Computer Science Dept., U. Rochester, expected 1982.

Cook, R.L. and K.E. Torrance. A reflectance model for computer graphics. *Computer Graphics (Proc. SIGGRAPH 1981) 15*, 3, 307-316, August 1981.

DeValois, R.L. and K.K. DeValois. Neural coding of color. In *Handbook of Perception*, Vol. 5, 117-166. E.C. Carterette and H. Friedman (Eds.). New York: Academic Press, 1975.

Gibson, J.J. *The Perception of the Physical World*. Boston: Houghton Mifflin, 1950.

Gilchrist, A.L. The perception of surface blacks and whites. *Scientific American 240*, 3, 112-124, March 1979. How lightness perception is influenced by surface slant perception.

Haber, R.N. and M. Hershenson. *The Psychology of Visual Perception.* New York: Holt, Rinehart and Winston, 1980. Standard (undergraduate?) text on visual perception.

Horn, B.K.P. Determining lightness from an image. *CGIP 3,* 277-299, 1974. Calculating the reflectance of surfaces given just their appearances in an image. Inversion of Laplacian operator, with various implementations and experiments. Part of a double whammy with Marr.

Horn, B.K.P. Exact reproduction of colored images. Working Paper 207, MIT AI Lab, December 1980. Good bibliography. Treats spectral distributions as infinite-dimensional vectors. Fundamental and pretty.

Horn, B.K.P and R.W. Sjoberg. Calculating the reflectance map. *Proc., DARPA Image Understanding Workshop,* November 1978. Mathematically intensive. Appears elsewhere. The reference for reflectance (achromatic, usually) calculations.

Judd, D.B. *Contributions to Color Science.* D.L. MacAdam (Ed.). NBS special publication 545, 1979. Collection of classic papers, including "theory" of surface color (from 1930) that may explain all Land's results.

Judd, D.B. and G. Wysecki. *Color in Business, Science, and Industry.* New York: Wiley, 1963. There's a later edition. Very nice introduction to "Color Science," or colorimetry theory and practice mixed with theoretical and practical aspects of psychology and psychophysics.

Land, E.H. The retinex theory of color vision. *Scientific American 237,* 6, 108-128, December 1977. Clear, fairly quantitative account of retinex theory (not the two-color projection results). More quantitative paper is McCann et al.

Land, E.H. Color vision and the natural image, parts I and II. *Proc., Nat. Acad. Sci. 45,* 1, January 1959). The basic two-color projection results. There's also a *Sci. American* version of these.

Marr, D. The computation of lightness by the primate retina. *Vision Research 14,* 1377-1399, 1974. How the primate retina might implement Horn's reflectance-extraction calculations.

McCann, J.J., S.P. McKee, and T.H. Taylor, Quantitative studies in retinex theory. *Vision Research 16,* 10, 1061-1069, 1976. Lots of raw results digested in Land's SA article. Many results reported (as usual) in Munsell color system.

Meyer, G. and D. Greenberg. Perceptual color spaces for computer graphics. *Proc., SIGGRAPH,* 1980. Good article on how to get your color monitor to give you colors in various other useful coordinate systems (such as Munsell). A bit terse in the basics--try Boynton or Judd and Wysecki, or Meyer's references. Meyer's thesis at Cornell should be a useful document, but I haven't seen it ... .

Mollon, J.D. Color vision. *Ann. Rev. Psychol. 33:* 41-85, 1982. Overview of recent work, with emphasis on color-opponent channels, short wave system anomalies, and fundamental cone responses.

Rock, I. In defense of unconscious inference. In *Stability and Constancy in Visual Perception.* W. Epstein (Ed.). New York: Wiley, 1977.

Rubin, Jim and W.A. Richards. Color vision and image intensities: When are changes material? AI Memo 631, MIT AI Lab, May 1981.

Sheppard, J.J., Jr. *Human Color Perception: A Critical Study of the Experimental Foundation.* American Elsevier, 1968. Good review of basic color experimental data, with warnings about unwarranted interpretations.

Wandell, B.A. Measurement of small color differences. *Psych. Rev. 89,* 3, 281-302, May 1982. Modern treatment of small perceptual differences. Interesting for its coverage of statistical model of color response. Lights are assigned values as (probability) functions on a 3-D space, not values as vectors in 3-D space.

Williams, D.R., J. Krauskopf, and D.W. Heeley. In search of the cardinal directions of color space. *J Opt. Soc. Am. 70,* 1574, 1980.

Zeki, S. The representation of colours in the cerebral cortex. *Nature 284,* 412-418, 1980.

Figure Captions

Figure 1:   Levels and nomenclature in color vision.

Figure 2:   Spectral reflectance of two pigments.

Figure 3:   The foveal cone spectral sensitivity function for a normal observer.

Figure 4:   Main components of the primate visual system.

Figure 5:   A quantitative representation of a specific opponent-color model. The curves R, G, and B represent the relative spectral sensitivites of the three kinds of cones. (Note that that these are weighted sensitivities for this model only -- in fact the unweighted B sensitivity is about two orders of magnitude below that shown). The R and G curves have been placed vertically to cross near 570 nm, and the B curve has been placed to cross the arithmetic sum of the R and G curves near 500 nm, so the y-b opponent system to occur near that wavelength. Opponent curves are calculated as shown by the labels attached to them in the bottom part of the figure.

Figure 6: A hypothetical connection net to compute opponent colors. The R, G, B cones at the top are the same as those at the bottom, repeated for clarity. The achromatic (nonopponent, luminance) pathway is activated by the sum of the R and G outputs,and the r-g pathway by their difference. The y-b pathway gets the difference between B and the luminance channel.

Figure 7: Plots of average firing rates of a large sample of cells from each of the six LGN cell types in response to flashes of monochromatic light. The top four are spectrally opponent cells that fire to some wavelengths and inhibit to others. The bottom two are spectrally nonopponent cells (from [DeValois and DeValois 1975]).

Figure 8: Wright's color-matching data.

Figure 9:   Tristimulus values of the 1931 CIE standard observer for the X, Y, Z primaries and equal energy white stimulus.

Figure 10: The (X, Y, Z) tristimulus space and the location of the (x, y) chromaticity diagram.

Figure 11:   The (x, y) chromaticity diagram for the CIE International (X, Y, Z) System.

Figure 12:   Color regions of the chromaticity diagram.

Figure 13:   Two color spaces, showing relation of opponent to HSV space.
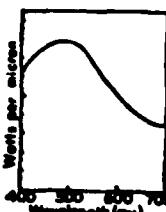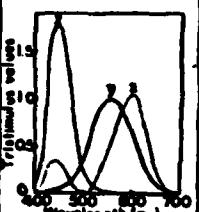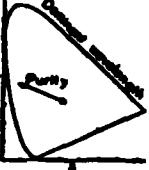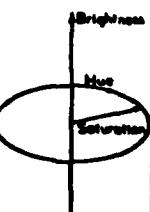
PHYSICS	PSYCHOPHYSICS	PSYCHOLOGY	
Visual Stimulus →	Light	← Visual Sensation →	Visual Perception
Radiant energy	Luminous energy		
Spectral composition →	Color	← Color sensation	
		Brightness	
		Hue	
		Saturation	
Characteristics of Radiant Energy	Characteristics of Light = Color	Attributes of Color Sensation	Corresponding Modes of Appearance
Radiant flux →	Luminous flux	← Brightness	Aperture (1-5)
Radiance →	Luminance		Illuminant (1-8)
Irradiance →	Illuminance		Illumination (1-3)
Radiant reflectance →	Luminous reflectance		Object modes: Surface (1-11)
Radiant transmittance →	Luminous transmittance		Volume (1-9)
Spectral distribution →	Chromaticity	← Chromaticness	Attributes of modes of appearance:
(Relative spectral composition, quality)	Dominant wavelength (or complementary)	Hue	1. Brightness (or lightness)
Radiant purity	Purity	Saturation	2. Hue
			3. Saturation
			4. Size
			5. Shape
			6. Location
			7. Flicker
			8. Sparkle
			9. Transparency
			10. Glossiness
			11. Luster

Figure 1

Green (blue ink on yellow ink)

Grass in sunlight
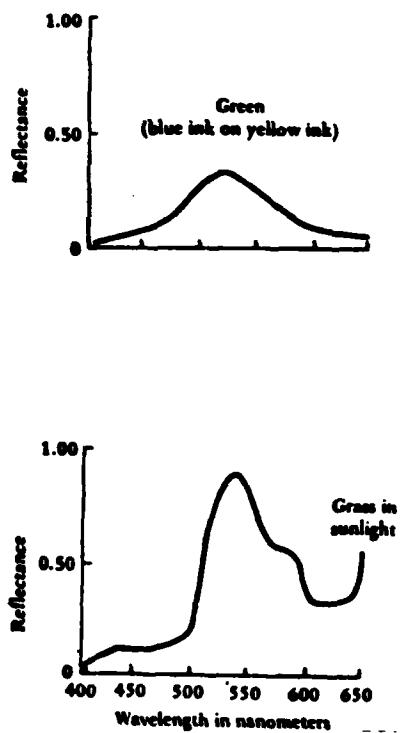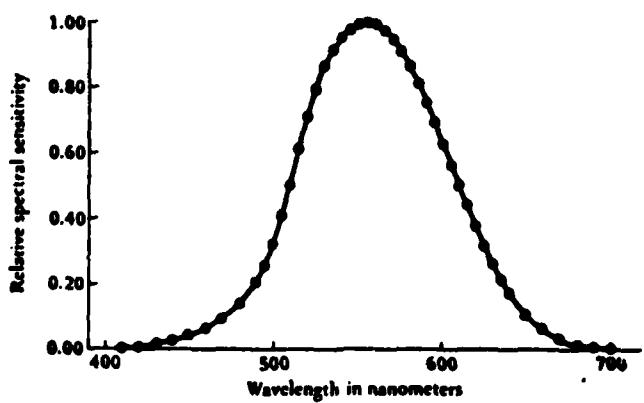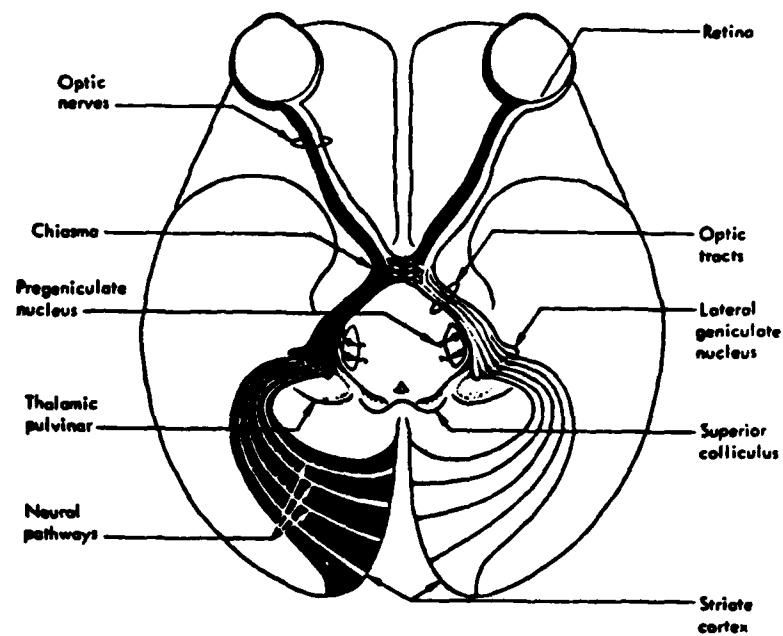
Figure 2

Figure 3

Figure 4
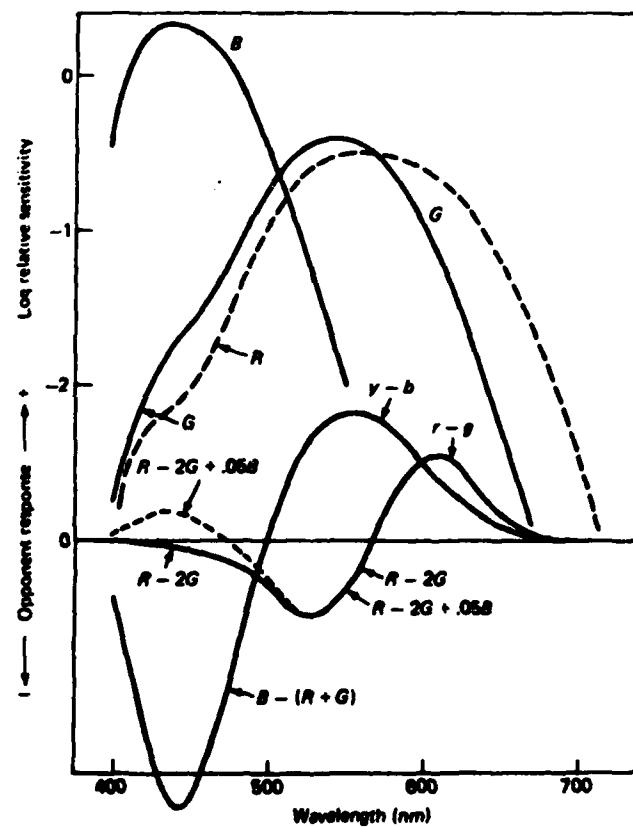

Figure 5
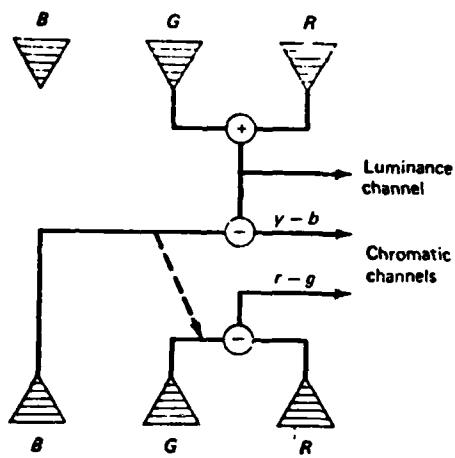
Figure 6
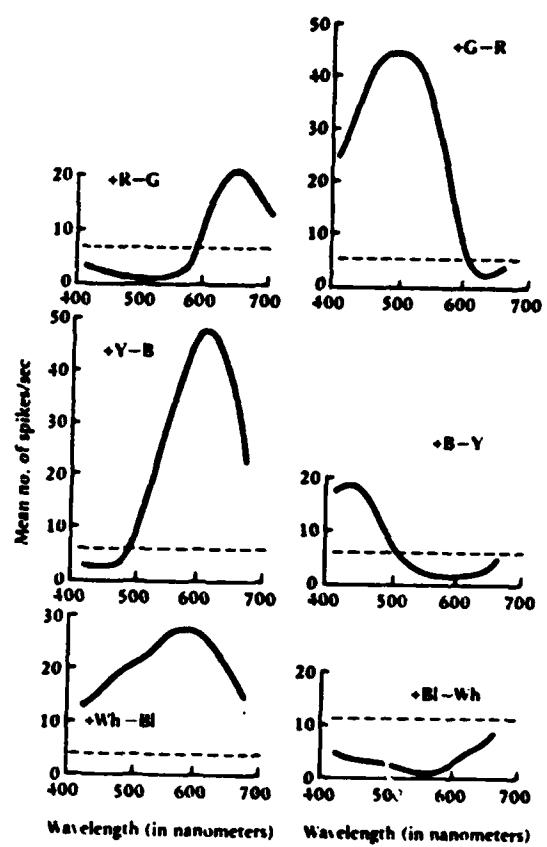
Figure 7
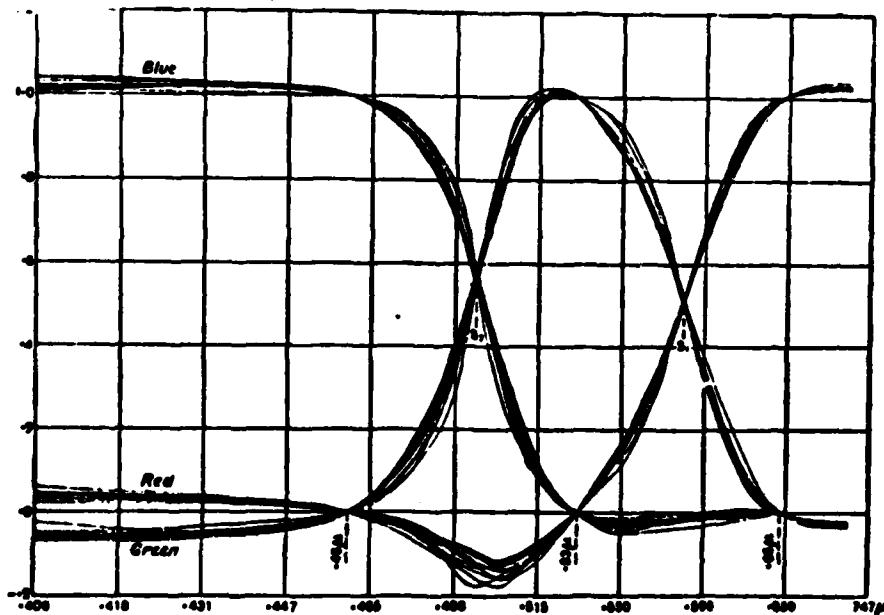
Figure 8
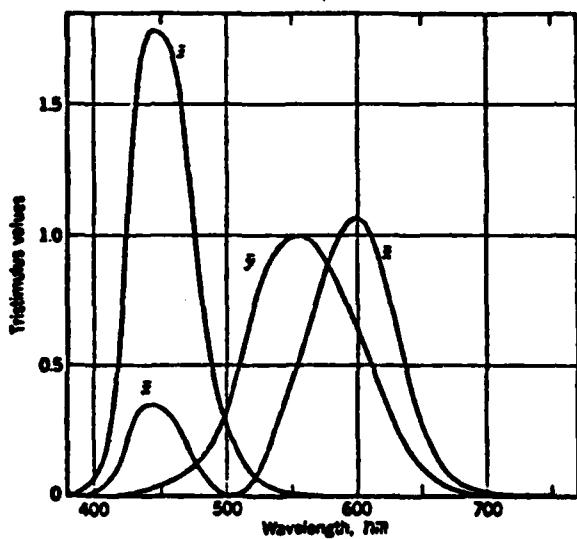
Figure 9
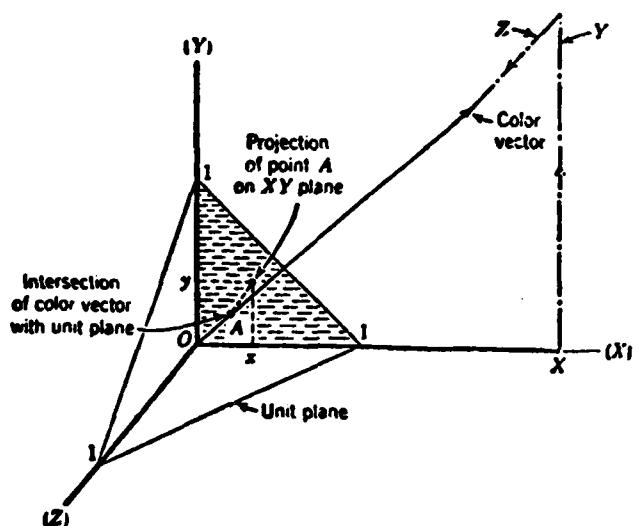
Figure 10
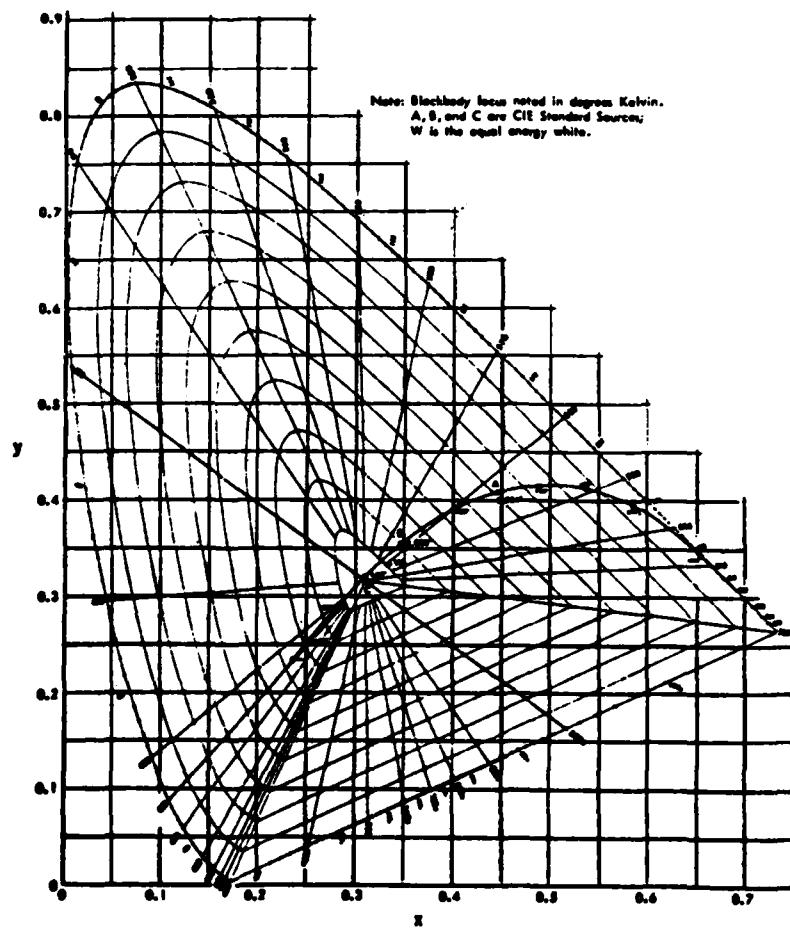
Figure 11
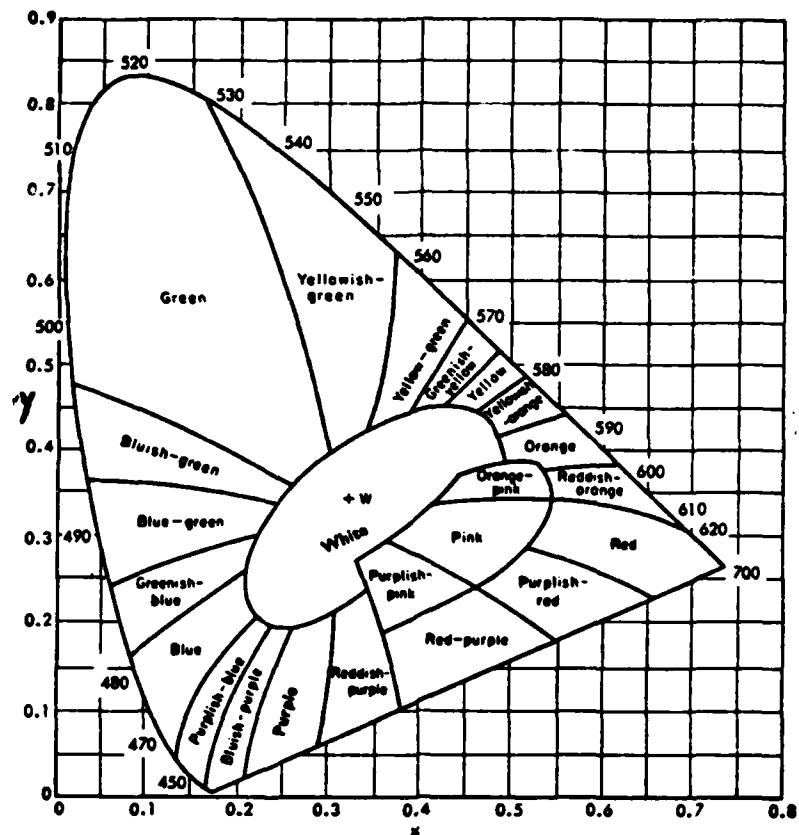
Figure 12
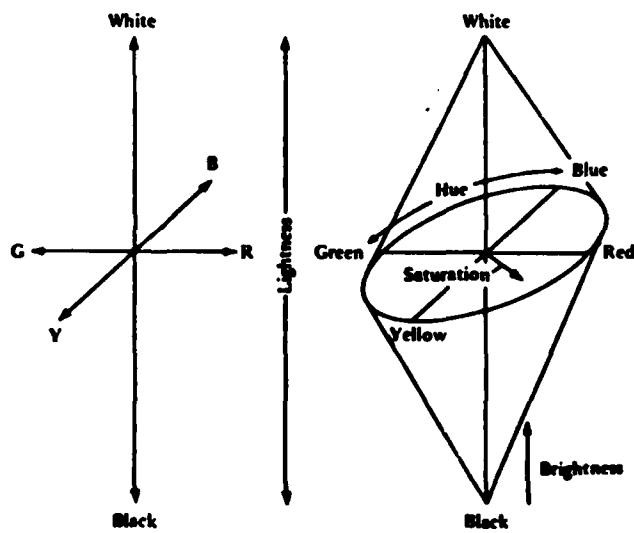
Figure 13

END

FILMED

9-83

DTIC